


CONNE~~X~~IONS



The Interoperability Report

September 1990

Volume 4, No. 9

ConneXions —

The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

TCP & TP4: Moving forward...	2
Components of OSI: Group Communication.....	8
INTEROP on July 4th.....	13
Alternative Book Review.....	14
Network Nomenclature.....	20
Upcoming Events.....	23

ConneXions is published monthly by Interop, Inc., 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. 415-941-3399. Fax: 415-949-1779.

Copyright © 1990 by Interop, Inc. Quotation with attribution encouraged.

ConneXions—The Interoperability Report and the ConneXions masthead are trademarks of Interop, Inc.

ISSN 0894-5926

From the Editor

Technically speaking, the OSI and Internet suite of protocols were both developed to solve the same problem: allow heterogeneous computer systems to interoperate. At the protocol level, similarities abound, reflecting the influence of one protocol suite upon the other. This is certainly the case with TCP and TP4 as discussed in this month's first article by two OSI and Internet veterans: Dave Piscitello and Lyman Chapin. However, beyond the technical similarities lies a large "culture gap"—Internet and OSI developers share little in the areas of perspective, working style, even terminology. This has led to many "protocol wars" and some fairly harsh criticism from both sides of the fence. The authors suggest that the energy used to fuel these fires could be better employed in a cooperative spirit as both protocol suites continue to evolve.

The Internet and OSI culture gap is perhaps best described in a number of "soap boxes" in Marshall Rose's *Open Book* which was published last year around this time. Some people have reacted with great praise, while others find the comments deeply offensive, and with this in mind I decided that it was time to start another series in ConneXions. I call it "Alternative Book Reviews." Realizing that any review is a very subjective thing, shaped by the experience and perspective of the reviewer, we will from time to time publish "the other point of view." First out is another review of the *Open Book* by Bryan Wood who has been involved in the OSI standardization process. For another opinion, see ConneXions, Volume 3, No. 11.

This month we have yet another installment in our series *Components of OSI*. Steve Benford describes the work being done to develop standards for Group Communication. E-mail can be much more than a replacement for traditional mail, and this area holds promising developments likely to affect the way we do business in the future.

INTEROP® 90 is now only one month away, and the Interop, Inc. headquarters is buzzing with activity as we prepare for this our annual event. Back in July, the engineers responsible for the INTEROP show network gathered at the San Jose Convention Center to design and build the network backbone. A brief report can be found on page 13. If you haven't registered for INTEROP 90, call us now at 1-800-INTEROP or 415-941-3399 to reserve your space.

There is some confusion among users and developers of networks as to what exactly is part of the Internet and what isn't. In an article starting on page 20, John Quarterman gives some suggestions for a more systematic nomenclature.

Finally some information about SIGCOMM '90 and the recently issued Internet Bibliography. See you at INTEROP!

TCP and TP4: Moving Forward

by

**David M. Piscitello, Bell Communications Research
and
A. Lyman Chapin, Data General Corporation**

In the beginning

The *Transmission Control Protocol* (TCP) and the *Open Systems Interconnection (OSI) Class 4 Transport Protocol* (TP4) are both designed to fill the architectural role of *end-to-end transport*—ensuring the reliable, sequenced, flow-controlled exchange of data between the two final endpoints of communication in an arbitrarily large and complex “internet” of interconnected networks. TCP was originally developed in 1973 as part of Robert Kahn’s “internet program” at DARPA’s Information Processing Techniques Office, and was formally published as a US Department of Defense Military Standard [1] in 1983. TP4 was developed as an OSI-compatible derivative of TCP, and was first published as an *International Organization for Standardization* (ISO) protocol standard [2] in 1986. Both are intended to operate in conjunction with an underlying connectionless internetwork protocol (IP).

Below IP, there is almost no difference (and therefore little argument) between the TCP/IP and OSI architectures. Above transport, the differences are considerable, and make fair comparison more difficult. TCP and TP4, however, are “almost the same.” Many people have tried to exploit this similarity to demonstrate the inherent superiority of the TCP/IP or OSI architecture, by comparing TCP to TP4 in a way that highlights the advantages of one and the deficiencies of the other. The authors suggest a much more productive exercise, in which the experience that has been gained by the developers and implementors of TCP accrues to the benefit of an evolving TP4.

Transport wars

It has become commonplace for computer industry trade journals to publish articles that position OSI and TCP/IP as contenders for multi-vendor network supremacy. Some of the articles read like a tout sheet for a heavyweight title fight: hostility, jealousy, competition, controversy, hype. In one corner, we have TCP/IP, representing over two decades of research and experimentation, and currently enjoying a tremendous market resurgence. In the other corner is OSI, a relative newcomer, whose reputation owes more to widespread international endorsement than to experience in the ring. The scene such articles depict plays like a protocol engineer’s version of “my dad can beat your dad.” “Transport wars” certainly appeals to a smaller audience than “Star Wars,” but the journalistic principle—that controversy is more newsworthy than harmony—is the same, and has produced many of the same results.

Like father, like son

Since imitation is considered to be the sincerest form of flattery, it is difficult to understand why TCP/IP and OSI are so often positioned as adversaries. The very existence of TP4 and its associated ISO standard Internetwork Protocol [3] is testimony to the success of TCP/IP. In *Internetworking with TCP/IP* [4], Douglas Comer points out that “TCP has been so popular that one of the International Standards Organization reliable stream protocols, TP4, has been derived from it.” The authors can also say with confidence that the same is true for ISO IP. And ISO’s recognition and imitation of TCP/IP has in some cases been reciprocated; the most recently developed TCP/IP intra-Autonomous System routing protocol, *OSPF* (Open Shortest Path First), began life as a derivative of the OSI Intermediate System to Intermediate System (router to router) intra-domain routing protocol.

Why fight?

If OSI is truly an offspring of TCP/IP, then why do TCP/IP lovers hate it so much? Among the many possible answers are:

- OSI may have evolved from TCP/IP, but in the process lots of good things were added (lost), so that OSI is technically superior (inferior) to TCP/IP;
- Because the architectures are so similar, products based on OSI and products based on TCP/IP compete for the same market; and
- TCP/IP (OSI) must be better than OSI (TCP/IP) because we developed it (the “not invented here” syndrome).

TCP/IP bigots typically contend that the OSI protocols carry too much excess baggage, and that many of the “really good things” about TCP/IP have been lost or perverted in the process of international standardization. “ISORMites” scorn the lack of upper-layer functionality in TCP/IP, and claim that since the “truly important” parts of TCP/IP have been incorporated into OSI, only stubborn refusal to change with the times keeps TCP/IP alive. (“ISORM” is an acronym for the “ISO Reference Model,” which is the basic specification of the OSI architecture).

However, it is not nearly as important to identify the origin of the controversy as it is to agree that the time has come to put it behind and move on. Left behind will be the shrinking cohort of hard-core TCP/IP partisans, who will continue to believe that if they ignore OSI long enough it will go away; and the similarly shrinking cohort of starry-eyed ISORMites, who refuse to learn any practical lessons about real networks from the TCP/IP people who have already been there. The authors will happily bid good riddance to both.

OSI is inevitable

To date, the deployment of large-scale OSI networks is quite limited. The imperatives driving the international acceptance of OSI, however, are considerable. OSI standards represent a powerful international alignment of computer vendor and telecommunication carrier interests, and represent, for many organizations, a major step towards market stability and predictability. Augmented by the emergence of government OSI procurement specifications (“GOSIPs”), the combined political effect of these two forces will ultimately make OSI the norm, rather than the exception, in many large contract bids. From a technical standpoint, the elaborate OSI upper layers architecture—which today is often maligned as unnecessarily complex and a performance-stifling burden—will be increasingly important as sophisticated distributed processing applications enter the computing mainstream.

Serious system designers and users alike have begun to realize that OSI’s supposed problems of performance and resource utilization are not due to some inherently threatening “OSI baggage” that robs systems of their processing potential, but rather to the fact that OSI implementations have only just begun to benefit from the experience and engineering expertise that comes (gradually) with implementation maturity. This process will accelerate dramatically as the expertise acquired by TCP/IP engineers is applied to the implementation and refinement of the corresponding OSI protocols—assuming, of course, that the TCP/IP engineers overcome their initial disdain for OSI, and that the OSI developers pay attention to what the TCP/IP engineers have accomplished.

TCP and TP4: Moving Forward (*continued*)

So is TCP/IP

In terms of deployment, the market presence of TCP/IP dwarfs that of OSI, and that presence is likely to continue growing at a rapid pace for many years, no matter how many GOSIPs are adopted by national governments. TCP/IP is widely available and understood. Most importantly, it works. And it has benefited enormously from 17 years of steady improvement in both the protocols and their implementations.

It is important to recognize that the strengths of OSI are for the most part also the strengths of TCP/IP; from a technical standpoint, the rationale for deployment of TCP/IP networks is no less compelling than the rationale for deployment of OSI networks. TCP/IP developers who recognize that there are useful lessons to be learned from the evolution of OSI (particularly in the upper layers) will continue to build and sell TCP/IP networks long after OSI has established itself in the marketplace.

So which is better?

Douglas Comer remarks in [4] that “TCP is a communication protocol, not a piece of software”. If this were in fact the basis for comparing TCP to other transport protocols, there would be no contest between TCP and TP4. A 1985 study performed jointly by the U.S. Defense Communications Agency and the National Academy of Sciences [5], concludes that TCP and TP4 are functionally equivalent and provide essentially similar services (see Table 1 below).

	TCP	TP4
Data transfer	streams	blocks
Flow control	bytes	segments
Error detection	checksum	checksum
Addressing	16 bit port #	variable tsap-id
Interrupts	urgent ptr	expedited data
Security	11 byte IP field	variable TP, IP fields
Precedence	3 bit IP field	16 bit TP, IP fields
Datagrams	UDP	ISO 8602
Disconnect	graceful	abrupt

Table 1: Comparison of TCP and TP4 Functions

There are bones to be picked here, to be sure. The TCP urgent data pointer provides a much more useful synchronization of urgent with normal data than TP4’s separate expedited data *Transport Protocol Data Unit* (TPDU). The single-byte granularity of the sequence space used by TCP, on the other hand, means that TCP sequence numbers wrap around their 2^{16} modulus much too quickly at very high network transmission speeds; TP4’s per-packet, rather than per-byte, sequencing avoids (or at least defers) this problem. In every case, however, the differences between TCP and TP4 can be eliminated by making relatively small changes to one protocol or the other. In fact, the groups responsible for the evolution of the two protocols have already begun to make some of these changes. ISO will soon approve an addendum to ISO 8073 which incorporates several improvements suggested by TCP implementation experience, including larger maximum TPDU sizes, finer maximum TPDU size granularity, selective acknowledgement, and request acknowledgement. An *Internet Engineering Task Force* (IETF) working group is currently studying how to expand, by 8 bits, the TCP sequence number field. The similarity between the two protocols is growing, and arguments about which protocol is “better” are therefore largely pointless.

It's not what you do, it's how you do it

Far more than differences between the protocols themselves, it is *implementation strategies* which separate TCP from TP4 today. The most widely-used TCP implementations (e.g., the TCP/IP that is bundled with Berkeley's 4.3BSD UNIX software) have undergone years of debugging, analysis, and fine-tuning of both the code and the underlying algorithms. Hundreds of talented engineers have contributed to the improvement of TCP implementations, with the result that TCP/IP-based systems today are generally reliable and adaptable, and perform well. The same cannot yet be said about OSI implementations, which have a much shorter track record. An extensive series of papers, most of them published as RFCs, documents the accumulated "smarts" of the TCP/IP world. TP4 implementors have only just begun to compile a written record of the techniques that can be used to improve the performance and flexibility of TP4 implementations. (An outstanding source of information describing a BSD UNIX implementation of TP4 can be found in [6]).

What have we learned?

TCP has undergone a number of significant changes since Vint Cerf and Robert Kahn first began working on it in 1973, including, in 1978, the divestiture of its original internetwork functions (which were captured in the separate IP protocol specification). Efforts to correct or improve the behavior of TCP implementations have left an interesting and instructive record, as we show in Table 2.

1973	TCP development begins at DARPA
1974 RFC 675	TCP version 1; First TCP implementations at Stanford, BBN, and UCL
1977 IEN 5	TCP version 2
1978 IEN 21	TCP version 3 (exeunt IP functions)
1978 IEN 40	TCP version 4
1981 RFC 793	TCP internet standard
1982 RFC 813	TCP window and acknowledgement strategy
1983 RFC 879	TCP maximum segment size and related topics
1983 MIL-STD-1777	US Dept. of Defense military standard for TCP
1984 RFC 896	TCP/IP congestion control
1988 RFC 1071	Computing the internet checksum
1988 RFC 1072	TCP extensions for long-delay paths
1989 RFC 1106	TCP big window and NAK options
1989 RFC 1110	A problem with the TCP big window option
1990 RFC 1122	Requirements for Internet Hosts— Communication Layers
1990 RFC 1144	Compressing TCP/IP headers for low-speed serial links
1990 RFC 1146	TCP alternate checksum options

Table 2: Seventeen Years of TCP

RFC 813 presents a good example of how enhancements to TCP based on implementation experience have been adopted and documented. In 1982, David Clark and others at MIT recognized and reported a TCP behavior they called *Silly Window Syndrome* (SWS). SWS arises when a natural break in the TCP data stream (a "push" point) causes a sender to divide the usable window it has calculated between two segments (the "usable window" is based on the receiver's offered window and the sender's knowledge of the number of outstanding unacknowledged bytes of data that are "in transit").

TCP and TP4: Moving Forward (*continued*)

When the receiver acknowledges one of these smaller segments, the sender is induced to calculate a correspondingly smaller usable window, and the transmission of a correspondingly smaller segment. As long as the sender keeps sending, and the receiver keeps acknowledging, there is no natural way for these usable window allocations to be recombined. Over the course of a long uninterrupted transfer, throughput will tend to decline, until the sender stops sending and the receiver's eventual acknowledgement of all of the data in transit resets the sender's usable window calculation to match the receiver's offered window.

Clark and company suggest a simple solution in RFC 813: the sender should refrain from sending any data at all if the calculated usable window is less than a certain threshold fraction of the offered window. When it receives a small segment, the receiver should discourage the transmission of further small segments by artificially reducing the offered window in its acknowledgement to the sender. These suggestions have been incorporated into most, if not all, current TCP implementations. [TP4 uses a window-based flow control scheme that is similar to TCP's, but is much less susceptible to SWS because its flow control units are packets, rather than individual bytes of data].

Applying TCP smarts to TP4

Delayed acknowledgement: The same RFC that describes SWS also explains how overly frequent acknowledgement incurs inordinately high processing overhead at the sender, and suggests that if the receiver expects more data to arrive soon (i.e., the push bit is not set) and there is no window update to send, the receiver should delay acknowledgement. This “smart” acknowledgement algorithm has been added to TP4 in ISO 8073, Addendum 4 (which is still undergoing review and balloting within ISO). The packet orientation of TP4 led to the adoption of an additional, related enhancement that permits *selective acknowledgement* of individual TP4 TPDU^s when the delayed arrival of one or more TPDU^s would otherwise prevent acknowledgement of all TPDU^s with higher sequence numbers.

MTU discovery: If a host knows the size of the largest packet that can be sent end-to-end to a particular destination without fragmentation by the underlying internetwork protocol, it can optimize its choice of segment length for TCP traffic to that destination. The long-established conventions for static MTU (*Maximum Transmission Unit*) determination were recorded by Jon Postel in 1983 in RFC 879. New work on dynamic discovery of MTU size has recently begun in a working group of the IETF. The same problem—how to maximize throughput by using the largest PDU that can be sent without incurring the overhead of internetwork fragmentation—afflicts TP4 networks as well. In the original specification of TP4, the set of allowable TPDU sizes was fixed, and consisted of the integral powers of two from 128 through 8192 bytes. To allow for the optimization of TPDU size with respect to the underlying internet, an enhancement to TP4 extends the range of TPDU sizes upward to approximately 2^{256} , and establishes a linear TPDU size granularity of 128 bytes (eliminating the arbitrary and rarely optimal power-of-two increments).

What about performance?

If functionality is not an issue, the obvious remaining question is “which protocol performs better?,” or more precisely, “do implementations of TP4 generally perform better than or worse than implementations of TCP?” According to at least one study, the answer depends on what you measure and how you measure it.

A group implemented both TCP and TP4 under UNIX 4.2BSD to “better understand the performance characteristics of TP4, and to compare TP4 with TCP” [6]. From the report, it appears that their methodology was unquestionably fair, and the results indicated that “except for the checksum algorithms, there were no protocol features which would lead to differences in implementation performance.” Their experiments also show that over both a 10 Mbps LAN and a general topology internet (the ARPANET), “performances of the two implementations were quite similar.” The results of the experiments in [6] suggest that the TP4 checksum imposes considerable processing overhead; in their LAN experiment, for example, TP4 running without the checksum was 20% better than TCP, but with checksum turned on, TP4 throughput was nearly 40% lower.

In conclusion

Some would argue that there are only two choices: “stick with TCP, ignore TP4, and pray that it goes away,” or “go with TP4, forget about TCP, and allow it to follow the path of all aging dinosaurs.” Neither is a possible nor likely scenario for the foreseeable future. We suggest that the intelligent alternative is to continue to work with TCP and to apply our accumulated TCP smarts to TP4. In so doing, we effectively hedge our bets. If TCP wins, a better TCP will have been developed; if TP4 wins, we won’t have to endure a future cursed with TP4 implementations that have not benefited from the expertise of the TCP/IP community. And while the standoff continues, you are all likely to have better transport protocols, irrespective of where your loyalties lie.

References

- [1] Feinler, E. J., Jacobsen, O. J., Stahl, M. K., & Ward, C. A. (Editors): *DDN Protocol Handbook*, pp 1-147–1-324, “Military Standard Transmission Control Protocol,” MIL-STD-1778.
- [2] ISO 8073-1986, International Organization for Standardization, “Data Communications—Connection-oriented Transport Protocol.”
- [3] ISO 8473-1987, International Organization for Standardization, “Data Communications—Protocols for Providing the Connectionless mode Network Service.”
- [4] Douglas E. Comer, “Internetworking With TCP/IP—Principles, Protocols, and Architecture,” Prentice-Hall, ISBN 0-13-470154-2.
- [5] “Transport Protocols for Department of Defense Data Networks: Report to the Department of Defense and the National Bureau of Standards,” National Academy Press, February 1985.
- [6] A. Bricker, L. Landweber, T. Lebeck, & M. Vernon, “ISO Transport Protocol Experiments,” MTR-86W00002, The MITRE Corporation.

Ed.: Views expressed in this article are those of the authors and do not necessarily represent views of Data General, or of Bellcore or any of its client companies.

DAVID M. PISCITELLO is a Member of the Technical Staff at Bellcore, where he is involved in the development of Switched Multi-megabit Data Service (SMDS). Before joining Bellcore in 1988, David was a manager and senior software engineer at Unisys (nee Burroughs), where he was engaged in product development of proprietary and open network architectures including BNA, TCP/IP, and OSI. David was a long-standing member of several ISO and ANSI committees and was responsible for the development of several OSI standards including ISO CLNP (ISO 8473).

A. LYMAN CHAPIN received his B.A. in Mathematics from Cornell University in 1973. Since 1977 he has worked at Data General Corporation as a designer and implementor of communications and networking systems, and is a Senior Consulting Engineer in the Systems Architecture Group. He is the chairman of the ANSI task group responsible for the development of OSI Network layer and Transport layer service and protocol standards (X3S3.3), vice-chairman of ACM SIGCOMM, and a member of the Internet Activities Board (IAB).

Components of OSI: Group Communication

by Steve Benford, University of Nottingham

Background

Group Communication has been introduced as a new work item for the standardisation of OSI Message Handling Systems, as specified in X.400/MOTIS [7, 8]. At the present time, the scope of standardisation is still under discussion and further input is being sought. This article explains why Group Communication is an important topic for standardisation and examines some of the key issues to be addressed.

Introduction

Perhaps the most significant advance introduced by OSI technology has been the specification of a variety of tools aiding the development of distributed applications. These include both specification and modelling tools such as the *Abstract Service Definition Conventions* and common supporting services such as the *Remote Operations Service* and *Directory Service*. Indeed, it could be argued that it is this infrastructure for application development which will prove to be the major advantage which eventually leads to the adoption of OSI over other network technologies.

Currently, only a few specific OSI applications, such as electronic mail, have been standardised. However, attention is now focusing on several new application areas including *Open Distributed Processing* and *Group Communication*. *Group Communication* has recently emerged as an area of great commercial interest and has been receiving considerable attention within the computing press under a variety of labels, including "Groupware" and "Computer Supported Cooperative Work." It has also been adopted as a new work item for the standardisation of OSI Message Handling Systems.

Limitations of existing services

One way to understand the need for *Group Communication* support is to examine the limitations of existing computer-based communication services.

Electronic mail is the most advanced of the computer-based communication services available today and therefore represents a suitable starting point for discussion. The adoption of the X.400/MOTIS standard for electronic mail is likely to result in a dramatic increase in the number of subscribers and, consequently, the volume of traffic. This communications explosion will exacerbate a number of problems which have already become apparent, even with todays relatively limited usage.

Information overload

First and foremost is the problem of "information overload" where users receive so many messages that potentially useful information becomes lost in a jungle of junk mail [1]. One particular way in which electronic mail systems encourage information overload is the provision of distribution list mechanisms enabling a message to be sent to hundreds of recipients with only minimal additional effort than is required to send to one. Furthermore, recipients are typically provided with minimal support for filtering out unwanted mail. This situation can be summarised by saying that electronic mail systems place too much power in the hands of senders at the expense of recipients.

The second major problem with electronic mail is that it is limited to inter-personal (one-to-one) communication. There is no real support for communication within groups of people (distribution lists provide minimal support for one-to-many communication but this may lead to more problems than it overcomes as described above).

For example, in an electronic mail system, each user effectively sees an incoming stream of messages arriving at his/her private mailbox. There is typically no support for sharing messages within a group (e.g., a team report) or for the long term, shared storage of information (e.g., an archive of past publications) [2].

Another major problem concerns the lack of *working context* for users. Mail systems support the transfer of messages between mailboxes. However, once a message has arrived, there is little support for its further processing (e.g., what should be done with it, how does it relate to other messages or to a specific task?). The underlying problem is that people typically work within a specific context at a given time, but this is not reflected by their mail system.

A variety of Group Communication applications have been developed which attempt to overcome these problems. Computer conferencing and bulletin board systems such as *COM*, *USENET News* and *EIES* focus on providing shared access to messages within a group. Other systems such as the *Information Lens* attempt to enhance message based communication by providing greater working context for users [3]. Although these systems introduce many important ideas for supporting group working, they also suffer from a number of limitations which severely reduce their future potential.

Scaling and interworking

Many existing Group Communication applications will not scale effectively. This may be due to the adoption of centralised rather than distributed architectures (e.g., *COM*) or may be due to lack of effective management functionality. Another critical limitation is the lack of interworking between the different applications. Users are likely to work with a variety of tools and will require a degree of interworking.

For example, a message produced in one application may be required as input to another. This lack of interworking may be largely due to the lack of a common infrastructure supporting the applications or to the use of specific local communication protocols. Furthermore, existing applications often provide limited functionality and are not readily extensible or configurable to a user's environment.

Many of the limitations of existing Group Communication applications might be removed if they were based on a common, flexible infrastructure. Such an infrastructure is provided within the OSI environment.

In summary, an examination of our current communication services reveals that, on the one hand, standardised electronic mail provides little support for group working and, on the other hand, existing Group Communication services lack the common infrastructure provided by the OSI standards. The solution to this problem would seem to lie in providing standard support for Group Communication within the OSI environment. However, before discussing how this might be achieved, it is necessary to take a closer look at the nature and scope of Group Communication in a more general sense.

What is Group Communication?

In its most general sense, Group communication refers to computer support for the communication processes which occur within groups of people who are working together to achieve some set of goals. This topic is also commonly referred to as "Computer Supported Cooperative Work" (CSCW), "Cooperative Work Technology" (CWT) or "Groupware."

OSI Group Communication (*continued*)

Group Communication has a different focus from other areas of distributed computing in that:

- Communication is considered within the context of a specific group or task. This is in contrast to the view of existing interpersonal communication services which are typically limited to considerations of how to pass information between individuals.
- The role of technology is to support communication between people. This differs from applications such as shared databases where, although groups of users may access the same information, no real communication occurs.

Examples

The scope of this new area is extremely broad. Groups are involved in a wide diversity of "activities," ranging from small-scale formalised office procedures to world-wide information services such as newspapers. The following list indicates just how diverse these activities might be:

- bulletin boards
- computer conferencing and news distribution
- face to face meetings:
 - brainstorming
 - committees
 - expert meetings (e.g., medical case conference)
 - joint editing
 - office and organisational procedures (e.g., travel application)
 - production of electronic newspapers, journals and periodicals
 - software design teams
 - voting, elections and opinion polls
 - seminars, lectures, tutorials
 - presentations
 - panel sessions (question and answer)
 - auctions
 - trading on the stock market
 - buying a house
 - producing international standards
 - managing large engineering projects

Just this short list is sufficient to demonstrate that Group Communication might impact on nearly every area of human working and that the resulting activities might differ in many ways. Examples of these differences might be:

- whether communication occurs in real-time or asynchronously.
- the activity specific information types that are involved.
- the number of participants.
- different patterns of communication (i.e., is there a well-defined sequence of communication acts?)
- whether there is a fixed goal or whether the activity is continuous.

Common requirements

In spite of the great differences between them, these diverse activities also have many common requirements. For example, the need for shared access to information and support for the management of activities within a given environment (e.g., establishing the participants involved in an activity). These common requirements suggest that it may be possible to identify a set of generic Group Communication services on top of which different activities can be built. It is these generic services which might be suitable for standardisation within the OSI environment.

Issues for standardising Group Communication

Group Communication has been adopted as a new work item within the X.400 standardisation framework. Due to the scale of the subject, the current focus of this work is on defining the overall scope of standardisation. The fundamental issue to be decided is whether standardisation should concentrate on one or two specific activities which are particularly relevant to messaging (e.g., bulletin boards) or whether it should address the generic support required for a broader range of activities. The former seems more attainable within a short timescale and belongs firmly in the messaging domain. The latter might result in a better long term solution, but would almost certainly have a major impact on other standards (e.g., the OSI Directory service).

The standards bodies are currently considering a number of possible scenarios for standardisation, ranging from minimal enhancement of Message Handling Systems, through the involvement of other standards such as the Directory to advanced scenarios to support specialised applications. In the author's opinion, the more advanced scenarios represent the best approach in the long term. In the interests of supporting the broadest possible range of future activities, standardisation should develop a general "framework" for OSI based Group Communication. This framework would provide a set of generic services for supporting Group Communication activities.

One previous Group Communication project within the OSI environment has suggested that a framework for Group Communication should contain a number of components [4]. These are briefly described below.

Abstract model for specifying activity support

There is a requirement for a modelling tool capable of specifying the common properties of activities in an abstract way. The model would define both the common information types involved in Group Communication and the abstract operations used to manipulate this information.

Integrated Services Architecture

Group Communication support is likely to involve a range of existing OSI services such as Message Handling, Directory and Management. These services need to be integrated and coordinated in a coherent manner. This requires the definition of an integrated service architecture which supports the functionality of the abstract model mentioned above. One approach already identified by previous work is the definition of a new OSI "Group Communication Service," responsible for the coordination of other OSI services on behalf of users involved in specific activities.

Specific service models

For each underlying OSI service which will support Group Communication, a specific service model has to be developed specifying exactly how the service is to be used. For example, in the case of the Directory service, new schema may have to be defined.

OSI Group Communication (*continued*)

One final comment: it is clear that, given the adoption of a suitably broad view of Group Communication, standardisation work will impact on a number of existing standards. This raises the question of whether the Message Handling Group is the ideal place for this work to occur. One possibility is that a new working group for Group Communication should be established with the aim of developing a separate standard over a suitably long timescale.

References

- [1] Jacob Palme, "You Have 134 Unread Mail! Do You Want To Read Them Now?," *Proceedings of the IFIP 6.5 Working Conference on Computer-Based Message Services*, North-Holland, 1984.
- [2] Hugh Smith, "The Requirements for Group Communication Services," *Research into Networks and Distributed Application: Proceedings of EUTECO 88*, North-Holland, 1988.
- [3] T.W. Malone, K.W. Grant, K-Y Iai, R. Rao and D. Rosenblitt, "Semi-Structured Messages are Surprisingly Useful for Computer-Supported Coordination," In *CSCW '86 Conference on Computer Supported Cooperative Work*, Austin, Texas, 1986.
- [4] Hugh Smith, Julian Onions and Steve Benford, "Distributed Group Communication—The AMIGO Information Model," Ellis Horwood, Chichester, England, 1989.
- [5] Julian Onions, "Components of OSI: The X.400 Message Handling System," *ConneXions*, Volume 3, No. 5, May 1989.
- [6] Steve Benford, "Components of OSI: The X.500 Directory Services," *ConneXions*, Volume 3, No. 6, June 1989.
- [7] ISO Standard 10021: "Information Processing systems—Text Communications—MOTIS," 1988.
- [8] CCITT Recommendation X.400: "Message Handling Systems: System Model—Service Elements," 1988.
- [9] ISO and CCITT: "Information Processing Systems—Open Systems Interconnection—The Directory," ISO 9594-1-8, CCITT X.500-X.521, 1988.

STEVE BENFORD is a lecturer at the University of Nottingham, England, where he is a member of the Communications Research Group in the Department of Computer Science. Since completing his PhD entitled "Research into the Design of Distributed Directory Services" in 1989 his primary research interests have moved into the domain of Computer Supported Cooperative Work. He is currently involved with the GRACE project at Nottingham, developing a Group Communication Service for the OSI environment. He is also actively involved in the new work item for the standardisation of group communication within ISO/IEC JTC 1/SC 18/WG 4 on messaging.

INTEROP 90 on The Fourth of July

A dozen or so dedicated networking professionals assembled at the San Jose Convention Center during the week of July 2nd to "pre-wire" the show network which is going to be used for INTEROP® 90.

18 miles of cable



The schedule for this year's INTEROP is such that the network crew only has about 8 hours to "hang the wires," thus as much advance design as possible is necessary. Armed with about 18 miles of cable, detailed floor (and ceiling) plans, and stacks of modular connectors, the INTEROP network crew built cable harnesses on the floor, ready to be hoisted into position in October.

All connections were tested for correct wiring, and one engineer even used a few finished sections of the cable to operate a small TCP/IP network between a couple of PCs for "proof of concept." At the end of the week, everything was carefully labeled, placed on spools and pallets and put into storage.


Head-start

"Luck would have it that the convention center was free for a whole week in July, allowing us to design and test the network backbone before the show," commented Peter de Vries, shownet manager, at Interop, Inc. "Having all the cable, this far in advance, and cut into the right-sized pieces, gives us a much needed head-start for the actual event," he added.

Demonstrations

The network will interconnect all participating vendors, and will have connections to the worldwide Internet and several other networks, nationally and internationally. A number of special technology demonstrations will take place including FDDI, ISDN, OSI GOSIP, PPP, SMDS, SNMP, 10Base-T, ONC/NFS and X Windows. Call us at 1-800-INTEROP or 415-941-3399 for the INTEROP 90 program.

"Connectorizing": John Romkey attaches connectors and tests each segment for correct wiring.



Part of the INTEROP engineering crew: Richard James, Karl Auerbach, Peter de Vries, Dave Bridgham, Nan Dorio, Mike Peter, Stev Knowles, Shirley Hunt, and Ron Natalie.



Alternative Book Review

The Open Book: A Practical Perspective on OSI, by Marshall T. Rose, Prentice-Hall, 1989, ISBN 0-13-643016-3. Previously reviewed in *ConneXions*, Volume 3, No. 11, November 1989.

Provocative

This is, and is intended to be, a provocative book and, partly in consequence, I have some significant criticisms of it. Nevertheless, I believe that its publication is an important event because the book represents (despite the provocations) a serious evaluation of the OSI standards from the perspective of an implementor with a significant involvement in the Internet community. Indeed, Marshall Rose may have a unique perspective since he has been involved as an implementor and user of the Internet protocols, and he has, also, been involved in the *ISO Development Environment* (ISODE), an implementation of OSI upper layer standards together with an emulation of the OSI Transport Service over the service provided by TCP/IP.

The book is very difficult to characterise, first because it is incomplete in its coverage both of OSI itself and of the comparison between the OSI and Internet protocols, and secondly because, interspersed with objective discussions of implementation issues, there are judgments based on purely qualitative and subjective criteria. In consequence, although the book is billed on the back cover as "one of the clearest expositions of OSI architecture and protocols" and proposed as a possible student text, it should be taken as one man's informed (but incomplete and sometimes partial) view of the OSI activity, and of the relationship between the OSI and Internet activities, and of the standards activity itself.

Organisation

The main element of the book is a description and evaluation of the OSI standards, with illustrations from the ISODE experience. A second important element is a discussion of the problems for the Internet community of coexistence with, and eventual transition to, the use of OSI standards; an issue of very great importance if we are to avoid a permanent division of the Information Technology (IT) community into two camps. Finally, and representing what I believe to be the least satisfactory (because least constructive) element of the book, scattered throughout and gathered in a self indulgent polemic at the end, are comments on the international standardisation process itself and on the comparative histories of Internet and OSI.

Introduction to Interconnection Standards

The main part of the book begins with an introduction to the problem of interworking and to the OSI and Internet activities. This starts with a very good statement of the need for interconnection standards; the fact that the value of information depends upon its availability to those who need it, hence the importance of being able to access or transfer information held in IT systems, and the essential inadequacy of access and transfer which is limited to a single make of equipment.

It goes on to outline the processes through which the OSI and Internet standards are agreed and, in doing so, introduces for the first time a contentious distinction between the OSI and Internet activities which surfaces periodically throughout the book; that the Internet work was able to focus on the solution to the (essentially technical) problem of interworking while the OSI work was (is) prevented from doing this by the way in which it is carried out. This distinction goes to the heart of what standardisation is about and I will return to it at the end of this review.

Flaws	<p>The outline of the OSI process has a number of important flaws. One of the most important is that the differences between the ISO/IEC and CCITT activities are not made clear; for example, the different focus of the major part of CCITT work (on enabling the operation of international telecommunication services), and the difference of status and usage between CCITT recommendations and ISO/IEC (and Internet) standards. It also fails to make clear that the functional standards process is not an after-the-event patch-up of problems but an essential part of defining the application of OSI standards to user needs.</p>
	<p>The introduction also gives an outline of the structure of standards defined by the <i>OSI Basic Reference Model</i> and of the OSI standards themselves. This is a point of major weakness for the book as a whole because there is only very limited discussion of the Reference Model and, in consequence, very important aspects of the OSI standards activity are not examined.</p>
Avoiding the Architects	<p>This may, perhaps, reflect the implementor's perspective of the book because there is, of course, a well known game connected with OSI called "Avoiding the Architects" in which "practical implementors" urge the need to get on with the practical business of designing protocols without getting tangled up in unreal "architectural issues" (and, yes, this does sound uncomfortably like the normal argument on any project under time pressure—and what projects are not?) for getting down to implementation with the bare minimum of overall system design.</p>
	<p>The fact is that the Reference Model has been vital in maintaining coherence within the set of OSI standards (making possible the potential richness of functionality that Marshall Rose welcomes in the upper layer standards), in providing a framework for longer term developments, and in specifying standards which placed minimum constraints on implementations. Since there is only limited discussion of the Reference Model there is an inadequate introduction to a number of important concepts and the use of many terms in the body of the book without sufficient (if any) explanation (e.g., layer, entity, service provider, address, application).</p>
End-to-end Services	<p>The body of the description of OSI standards is in two parts. The first part deals with what are termed the "End-to-end Services" (provided by the lower four layers of the Reference Model), and the second, and larger, part deals with the "Application Services" (provided by the upper three layers).</p>
	<p>The treatment of the lower layers is only a partial one. It concentrates on the functionality of the <i>Network</i> and <i>Transport</i> Layers, referring elsewhere for detailed discussion of specific transmission media. The section provides a good introduction to the basic requirements to be addressed by the Network and Transport Layer standards, to the Network Service and the protocol structures in the Network Layer and to the Transport Service and Protocol. The principle weaknesses in the description are really matters of emphasis. In particular, it fails to bring out the critical importance of the concept of the Network service as the basis for providing transparency of underlying subnetworks and for interlinking them, and of the independence of the Network address from underlying addressing and routing mechanisms.</p>

Alternative Book Review (*continued*)

Not surprisingly, a major part of this section is concerned with an evaluation of the Network and Transport standards, in particular with reference to the TCP/IP standards used within the Internet community. Unfortunately, the value of the reference to the TCP/IP standards is limited since there is no description of these standards and hence the validity of some of the comments must be taken on trust. Some of the specific issues raised seem to need rather more examination, for example:

- *The complexity of the internal organisation of the Network Layer*—there is an implication in the discussion that the TCP/IP structure is inherently simpler, but this seems likely to be true only in the sense that the TCP/IP structure makes a choice of a single (connectionless-mode) Network service; providing support for that single Network service over both connectionless-mode and connection-mode subnetworks is not itself simplified by the choice of the TCP/IP protocols.
- *The possibility of choosing a single mode (connectionless-mode or connection-mode) of Network Service to support the connection-mode Transport service*—there is an implication that someone, somehow, should have imposed a choice on the standardisation process—but this reflects a failure to accept the essential nature of an international standardisation process, namely that agreement cannot be imposed because there is no one to impose it. The failure lies, perhaps, not so much in not making a choice as in not accepting that a choice could not be made and that both modes should be fully accommodated.
- *The desirability of choosing the connectionless-mode Network Service*—there is also a strong view expressed that the choice of a single mode of Network service is desirable, and an implication that the TCP/IP approach (i.e., a connectionless-mode Network Service) had clearly established its capability to answer all needs.

This ignores the fact that, despite their comparative maturity, the TCP/IP protocols have only been widely used in a limited number of areas of IT application, that their operation has not been without problems, and that their use raises problems of security, accounting and management in multi-domain environments. From a technical point of view the choice is still an open issue and likely to remain so for some time to come. Thus, from the user's point of view, it remains desirable to be able to make a choice between connection-mode or connectionless-mode Network operation based on the system needs and the characteristics of the transmission facilities to be used.

Application Services

The description of *Application Services* (addressing the standards for the upper layers) is much more substantial than that of the lower layers and represents the major part of the book. This reflects, perhaps, the focus of the ISO/IEC work and the area of most direct experience for Marshall Rose.

This section provides good, clear and detailed descriptions of the *Session Layer* standards, the *Presentation Layer* and ASN.1 standards, *Association Control*, *RTSE* (Reliable Transfer Service Element) and *ROSE* (Remote Operations Service Element) with illustrations of implementation issues taken from the ISO/IEC work.

A particularly valuable aspect of these descriptions is the recognition that the objective of the OSI work has been to establish a framework for the evolution of standards for application interworking which can be used together efficiently and effectively.

In particular, there is recognition of the central importance of the Presentation and ASN.1 standards in allowing the decoupling of information exchange between systems from the way in which that information is held within systems. The section also provides descriptions of the *Directory*, *Message Handling* and *FTAM* (File Transfer, Access and Management) standard but these are comparatively high level, functional descriptions, for example, in the case of FTAM, the *Reliable Service* is not considered.

Missing Standards

The section seems to be limited to standards which were fully agreed (if not yet formally published) at the time of writing and does not consider future developments. Thus there is no treatment of the *Commitment, Concurrency and Recovery* (CCR) standard, or of the *Virtual Terminal*, the *Job Transfer and Manipulation* (JTM) or the *Transaction Processing* standards. Furthermore, there is no discussion of the potential development of Application Layer standards by groups with particular needs, for example the development of the *Manufacturing Message Service* standard. Finally, there is no discussion of the work on OSI Management standards (although this work is referenced in the final section of the book which discusses and compares the processes followed in OSI and Internet standardisation processes).

The section is also weakened by the incomplete discussion of the OSI Reference Model and of the complementary *Application Layer Structure* standard so that there is not an adequate explanation of some concepts and terminology which are key to the understanding of the structuring and objectives of standards in the Application Layer. One important effect of this is that it is not made clear how the concept of application service element applies to the standards for Directory, Message Handling, and FTAM—for example, it is not made clear that the Message Handling standard involves specifications for a number of ASEs.

Transition and coexistence

A very significant part of the book is concerned with the need:

- to consider the mechanisms for coexistence in the short- to medium-term of networking based on the Internet standards and networking based on OSI standards—because of the existing (and continuing) investment on applications using Internet standards; and
- to consider strategies for transition from use Internet standards to use of OSI standards because of the support for the OSI programme and the increasing commitment of suppliers.

This attention is to be welcomed. It is clearly desirable for the evolution of networking that there should be a single framework of standards within which resources and effort can be concentrated. At the same time it is clearly undesirable and impractical to expect that existing systems should be replaced before the end of their useful life. This section is valuable both for its consideration of the scale of the problems to be faced and for its detailed examination of the mechanisms which are possible and of the strategies that can be used.

Alternative Book Review (*continued*)

Of particular interest is the concept drawn from the ISO/IEC experience that, as part of a transition strategy, already established networks with a TCP/IP based infrastructure could utilise OSI upper layer standards and establish an application environment based on those standards before it becomes practicable to utilise OSI lower layer standards. Given the widespread use of TCP/IP protocols in UNIX environments, it is also significant to see that part of the DoD transition strategy links a POSIX based application environment with communication functions based on OSI standards.

The standardisation process

The final part of the book together with (clearly identified) interpellations throughout the text, provides reflections on the standards process and a comparison of the OSI and Internet activities.

I found this to be the least satisfactory element of the Open Book because the manner in which it is carried out, apart from being gratuitously offensive to a large number of people, obscures and trivialises important issues about the nature, objectives and possible procedures for the standards process. The author may be entitled to make judgments about representatives to the standards process but since I know the people involved well, I object very strongly to the implication that a large proportion of the experts involved in the OSI activity are "nitwits" attending the meetings only to travel and drink.

It is, in any case, difficult to square the process described by Marshall Rose with the production of anything of value at all, let alone a substantial body of standards which, however flawed, are being incorporated by suppliers in mainline products. It would be an interesting sociological study to examine why the combination of networking protocols and standardisation should provoke such violent and unprofessional antipathies!

Problems with the ALS

Out of a number of specific issues raised I will comment on two. The first is the "rubbishing" of the work on the Application Layer Structure on the grounds of its lack of completeness and comprehensibility. The limitations of the ALS standard are well recognised. What is not made clear in these comments is that there has been no alternative proposal, and that the problems reflect the difficulties that have emerged in the attempt to develop a framework to give coherence to Application Layer standards in the long term. Would it have been better to do nothing?—bearing in mind that doing nothing would not have meant that no views of the ALS existed but that everyone had their own (unstated) view.

Myths of formalism

The second concerns the so called "myths of formalism"—formal specifications and conformance testing. The views expressed reflect a fundamentally different perspective on the nature of standards from that underlying the OSI work. It is perceived as a requirement for OSI standards (as with standards in other areas of technology) that the specification should be expressed in a way that places the minimum constraint on implementation and that makes clear the requirements for conformance to the standard.

Thus, standards are not specifications of implementations (even if they are defined in terms of a state machine), they are specifications of the expected behaviour of implementations.

It follows that the specification of conformance and the ability to test for conformance are integral to this concept of a standard. It is, of course, recognised that there are serious difficulties in fully working out such a concept but if, to achieve interworking, the implementation of standards for communication had to remain a black art then we would have a lot to fear from the spread of networking.

However, the overriding issue raised by this aspect of the Open Book relates to the processes needed to achieve effective and timely international standards—where effective means that standards both provide the functions for which they are developed and are supported by suppliers of systems and software, and timely means that they are available before there is a large investment in alternatives. From this point of view the Internet experience is very valuable because of the way in which progress was tied into experience and there are two aspects of the experience which it is particularly worthwhile to consider in relation to the OSI activity.


First, there had to be a “bootstrap” process for the Internet standards—standards had to be produced in advance of widespread implementation before there was an environment for prototyping and testing. The Open Book makes clear that we are in a similar “bootstrap” process now for OSI standards—and suggests a number of ways in which the Internet experience can ease and speed that process.

Difference in context

Secondly, the Internet process has been dependent on a “patron” and the control structure which that patron sponsored. Such a “patron” could not exist for international standardisation and a similar control structure could not be established. It follows that international standards must be established on the basis of agreement—they cannot be imposed. From this point of view there is not, as is suggested, a difference of technical focus between the OSI activity and the Internet activity. The vital difference is not of technical objectives but of the context within which those objectives are pursued. If we want open networking then we have no alternative but to accept the disciplines of international development and international agreement and to make them work as effectively as possible. —*Bryan Wood*

BRYAN WOOD is chairman of IST/21, the UK committee contributing to the work of the international JTC1/SC21 committee and Principal Consultant with the SEMA Group. Mr. Wood has reviewed this book in his personal capacity.

This book review is reprinted with permission from *OSN: The Open Systems Newsletter*, December 1989. OSN is published by Technology Appraisals Ltd, based in the United Kingdom. For an *entirely different* point of view on this book, see *ConneXions*, Volume 3, No. 11, November 1989. —*Ed.*



INTEROP® 90

October 8–12, 1990 • San Jose, CA

Network Nomenclature

by John S. Quarterman, Texas Internet Consulting

Introduction

There is some confusion among users and developers of networks as to what exactly is part of the Internet and what isn't. This article gives some suggestions for a more systematic nomenclature.

Basic categories

This nomenclature involves six basic categories, with each category including more networks than the previous one. There are two basic characteristics that are used to distinguish the categories:

- *Type of interactive connectivity*: Four of the categories involve interactive connectivity to the Internet, and are distinguished according to whether that connectivity is by use of IP or not, and whether it involves logging in on an intermediate gateway or not.
- *Commercial use*: Other networks are more loosely connected by high level protocols such as mail or news; the two categories proposed for them are distinguished by economic and political issues.

Each of these two distinguishing characteristics and six basic categories is discussed in turn below.

Four terms for networks with *interactive connectivity*, defined from the point of view of the Internet are summarised in Figure 1.

Each host on the Internet has end-to-end interactive IP connectivity from the user's machine to NSFNET. There are actually other agency networks that could serve equally well as a distinguishing central network. NSFNET is used here because it is the most widely known.

	IP		Non-IP	
	direct	gateway login	protocol translation	gateway login
Internet	•			
Catenet	•	•		
Transnet	•	•	•	
Outernet	•	•	•	•
Examples:	NSFNET	HP Internet	STARLAN	Easynet

Figure 1. Interactive connectivity.

The Internet alone is huge, including not only all the NSFNET regionals and associated campus networks, and numerous other networks in the United States, but also many other networks in other parts of the world, such as NORDUnet, InterEUnet, JUNET, and large parts of ACSnet.

Interactive IP connectivity by means of intermediate login to a gateway on the Internet distinguishes the *Catenet*. This includes the SGI and HP Internet situations.

I am taking some liberties with the historical use of this now-disused term. The first paragraph and the first reference of "The Catenet Model for Internetworking" by Vint Cerf, DARPA/IPTO, IEN 48, July 1978, read:

The term "catenet" was introduced by L. Pouzin in 1974 in his early paper on packet network interconnection [1]. The U.S. DARPA research project on this subject has adopted the term to mean roughly "the collection of packet networks which are connected together."

This is, however, not a sufficiently explicit definition to determine, for instance, whether a new network is in conformance with the rules for network interconnection which make the catenet function as confederation of co-operating networks. This paper attempts to define the objectives and limitations of the ARPA-internetworking project and to make explicit the catenet model on which the internetworking strategy is based.

Cerf also remarked in the mailing list `ietf@venera.isi.edu` on 19 February 1990 that:

The Catenet term was introduced by Louis Pouzin, then at IRIA, now called INRIA (Institut National de Recherche en Informatique et Automatique). He meant what the term "Internet" has come to mean (everything you can reach with IP, assuming you are on the same system as the backbones—NSFNET, MILNET, NSINET...)

Transnet	Interactive connectivity from a non-IP network to the Internet by transport or higher layer protocol translation distinguishes <i>Transnet</i> . This happens from the AT&T Bell Laboratories STARLAN network.
Outernet	Interactive connectivity from a non-IP network to the Internet by means of intermediate gateway login distinguishes <i>Outernet</i> . This includes DEC's Easynet, the XEROX Internet, and JANET

	<i>Internet connectivity</i>	<i>dedicated links</i>	<i>dialup networks</i>
Worldnet	•	•	•
The Matrix	•	•	•
<i>Examples:</i>	Internet Catenet Transnet Outernet	BITNET HEPNET XEROX I. SPAN	UUCP USENET FidoNet others

Figure 2. More inclusive terms.

Many networks do not have interactive connectivity. There is less point in defining groups of these from the point of view of the Internet, or from any central point of view. The glue here is electronic mail or news, although other protocols and services, such as batch FTP, may also be supported. Two categories are presented here, and are shown in Figure 2 with examples.

WorldNet The set of all research, academic, or cooperative networks that exchange mail or news is *WorldNet*. This includes all four of the terms already introduced (the Internet, the Catenet, Transnet, Outernet), plus BITNET, USENET, UUCP, CDNnet, FidoNet, and the Kermit link to Antarctica. The name WorldNet was apparently invented on the HUMAN-NETS list in the seventies (if anybody has more definite information, I'm very interested) and was then meant to apply to all computer networks worldwide. However, inasmuch as it is still used at all, it has come to have a more specific application that excludes for-profit commercial systems.

The Matrix The set of all systems that exchange electronic mail or news (batch Computer Mediated Conferencing), regardless of funding, administration, or purpose, is the *Matrix*. This includes everything on WorldNet, plus commercial systems such as CompuServe, the Source, BIX, and Minitel. It does not include isolated machines (such as single-PC bulletin board systems that allow only user dialups) or networks that are not connected to other networks.

continued on next page

Network Nomenclature (*continued*)

The term, the matrix, was apparently invented by science fiction writer William Gibson, and appears in his books such as *Neuromancer* to indicate a worldwide network used by almost everyone and involving very high speeds, sophisticated graphics, and elaborate conventions for iconic representation. This network doesn't yet exist in the form he describes. However, he acknowledges that the source of the idea is Vernor Vinge's story *True Names*, which explicitly mentions the ARPANET as one of the predecessor networks. I use the name the Matrix for the current worldwide network, because it is the precursor of what Gibson describes. It's also simply a fitting name.

The distinction between WorldNet and the Matrix is somewhat controversial. Some think that WorldNet should include all connected networks. However, historically, the academic and research community has deliberately isolated itself from the commercial world by limiting access to academic and research networks to non-commercial uses and by holding workshops and conferences at which industry representatives did not feel altogether welcome. Nonetheless, this distinction is the most arbitrary of all those given here, and is becoming more blurred daily, now that networks like CompuServe and MCI Mail are interconnected with the Internet and services such as those of ClariNet and Anterior Technology are carried over the Internet.

Related terms

Similar categories and tables could be built starting with, e.g., JANET, and continuing with Transnet, Outernet, WorldNet, and the Matrix. The definitions of Transnet and Outernet would have "JANET" substituted for "the Internet" and "Coloured Book" substituted for "IP." Outernet would be the same, but the Internet would be in the outer part of it from that point of view. WorldNet and the Matrix would be the same.

There are some interesting subsets of networks within the Matrix that can be distinguished according to the use of certain high-level protocols. For example, USENET isn't entirely a dialup network. Much of it is carried over NSFNET; some is carried over ACSnet, etc. Its distinguishing characteristic is that all its hosts carry USENET news. This means they understand the USENET article format defined in RFC 1036 (well, except for those that run *notesfiles*) and that they carry a certain minimal subset of newsgroups.

The set of all hosts that understand Internet *Domain Name System* addresses forms a sort of logical network that incorporates (most of) the Internet, BITNET, and EUnet, and parts of UUCP and other networks. This is often confused with the Internet itself. Perhaps a better name for it would be the *Internet DNS Mail System*, or perhaps the *Internet Mail System*, for short.

References

- [1] Pouzin, L., "A Proposal for Interconnecting Packet Switching Networks," Proc. EUROCOMP, Bronel University, May 1974.

JOHN S. QUARTERMAN is a partner in Texas Internet Consulting of Austin, which specializes in UNIX systems programming, UNIX and network-related standards, and local area network design, installation, and management. He has previously worked for the University of Texas and BBN, and graduated from Harvard (A.B. 1977). He is a coauthor of the definitive book on 4.3BSD UNIX and the author of the most comprehensive book about actual computer networks, *The Matrix*. Quarterman is also involved with a new company, Matrix Information and Directory Services, Inc., whose purpose is to provide current and comprehensive information about actual networks. He can be reached as jsq@longway.tic.com.

Upcoming Events

Networking in the year 2000

The *ACM SIGCOMM '90 SYMPOSIUM: Communications Architectures and Protocols* will be held in Philadelphia, Pennsylvania, September 24–27, 1990. The symposium is the annual conference of ACM SIGCOMM and provides an international forum for the presentation and discussion of communication network applications and technologies, as well as recent advances and proposals on communication architectures, protocols, algorithms, and performance models. This year the conference particularly encouraged papers on the topic of networking in the year 2000 (making networks very fast and very big).

Tutorials

The conference offers two tutorials on September 24. Tutorial #1 on *Protocols for High Speed Networks* will be taught jointly by Harry Rudin and Van Jacobson. Tutorial #2, on *Object-Oriented Network Management and Control* will be taught jointly by Aurel Lazar and Mark W. Sylor.

Technical program

The technical program, September 25–27, includes 31 papers (selected from over 100 papers submitted) presented in a single track format. The ten session topics are:

- Congestion Control
- Applications & Distributed Systems
- MANs and WANs
- Multimedia Protocols & Protocol Testing
- Gigabit Protocols
- High-Speed Switching
- Routing & Flow Control
- Protocol Design
- LAN Issues
- Routing

ACM SIGCOMM Award

The keynote speaker for the Technical Program will be the winner of the annual *ACM SIGCOMM Award* for lifetime contributions to the field of computer communication. The award winner will be announced at the conference.

For more information about and registration material for the conference, see the advertisements in the July issues of *Computer Communication Review* and *Communications of the ACM*, and the August issue of *IEEE Communications Magazine*. You can also contact the SIGCOMM '90 information office by sending e-mail to sigcomm90@cis.upenn.edu or by calling 215-898-0016.

Internet Bibliography now available

A new Internet Draft is available from the on-line *internet-drafts* directories called "Where to Start—A Bibliography of General Internetworking Information."

The filename is `draft-ietf-userdoc-bibliography-00.txt`

Information sources

The intent of this bibliography is to offer a representative collection of resources of information that will help the reader become familiar with the concepts of internetworking. It is meant to be a starting place for further research. There are references to other sources of information for those users wishing to pursue, in greater depth, the issues and complexities of the current networking environment.

Getting the draft

This draft can be obtained via anonymous FTP from the *internet-drafts* directories with the filename above. Internet-Drafts are available from the following hosts: nnsa.nsf.net, nic.ddn.mil, munnari.oz.au, and nic.nordu.net.

—Greg Vaudreuil

CONNEX**IONS**
 480 San Antonio Road
 Suite 100
 Mountain View, CA 94040
 415-941-3399
 FAX: 415-949-1779

Bulk Rate
 U.S. POSTAGE
 P A I D
 SAN JOSE, CA
 PERMIT NO. 1

ADDRESS CORRECTION
 REQUESTED

CONNEX**IONS**

EDITOR and PUBLISHER **Ole J. Jacobsen**

EDITORIAL ADVISORY BOARD **Dr. Vinton G. Cerf, Vice President,
 Corporation for National Research Initiatives.**

**A. Lyman Chapin, Senior Consulting Engineer,
 Data General Corporation.**

**Dr. David D. Clark, Senior Research Scientist,
 Massachusetts Institute of Technology.**

**Dr. David L. Mills, Professor,
 University of Delaware.**

**Dr. Jonathan B. Postel, Communications Division Director,
 University of Southern California,
 Information Sciences Institute.**

Subscribe to CONNEX**IONS**

U.S./Canada \$125. for 12 issues/year \$225. for 24 issues/two years \$300. for 36 issues/three years

International \$ 50. additional per year **(Please apply to all of the above.)**

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

Check enclosed (in U.S. dollars made payable to **CONNE**X**IONS**).

Charge my Visa MasterCard Am Ex Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to: **CONNE**X**IONS**

Back issues available upon request \$15./each
 Volume discounts available upon request

480 San Antonio Road Suite 100
 Mountain View, CA 94040
 415-941-3399 FAX: 415-949-1779